arXiv:1408.2356v1 [gr-qc] 11 Aug 2014

Testing general relativity with compact coalescing binaries: comparing exact and predictive methods to compute the Bayes factor

Walter Del Pozzo ‡, Katherine Grover, Ilya Mandel, Alberto Vecchio

School of Physics and Astronomy, University of Birmingham, Edgbaston, Birmingham B15 2TT, United Kingdom

Abstract. The second generation of gravitational-wave detectors is scheduled to start operations in 2015. Gravitational-wave signatures of compact binary coalescences could be used to accurately test the strong-field dynamical predictions of general relativity. Computationally expensive data analysis pipelines, including TIGER, have been developed to carry out such tests. As a means to cheaply assess whether a particular deviation from general relativity can be detected, Cornish et al. [11] and Vallisneri [15] recently proposed an approximate scheme to compute the Bayes factor between a general-relativity gravitational-wave model and a model representing a class of alternative theories of gravity parametrised by one additional parameter. This approximate scheme is based on only two easy-to-compute quantities: the signal-to-noise ratio of the signal and the fitting factor between the signal and the manifold of possible waveforms within general relativity.

In this work, we compare the prediction from the approximate formula against an exact numerical calculation of the Bayes factor using the `lalinference` library. We find that, using frequency-domain waveforms, the approximate scheme predicts exact results with good accuracy, providing the correct scaling with the signal-to-noise ratio at a fitting factor value of 0.992 and the correct scaling with the fitting factor at a signal-to-noise ratio of 20, down to a fitting factor of ~ 0.9. We extend the framework for the approximate calculation of the Bayes factor which significantly increases its range of validity, at least to fitting factors of ~ 0.7 or higher.

‡ wdp@star.sr.bham.ac.uk

1. Introduction

The upgraded versions of the ground-based gravitational wave detectors LIGO [1, 2] and Virgo [3, 4, 5, 6] are expected to detect gravitational-wave signals from the coalescence of compact binary systems. The prospect of frequent detections, with expected rates between one per few years and a few hundred per year [7], promises to yield a variety of scientific discoveries. Among these, the possibility of testing the strong field dynamics of general relativity (GR) has received increasing attention (e.g., [10, 11, 12, 14, 13, 9]). In fact, during the latest phase of the inspiral, typical orbital velocities are an appreciable fraction of the speed of light ($v/c \sim 0.4$); following merger, the *compactness* $GM/(Rc^2)$ of the newly formed black hole that is undergoing quasinormal ringing is close to 1. By comparison, the orbital velocity of the double pulsar J0737-3039 is $O(10^{-3}c)$ and its compactness is $\sim 10^{-6}$ [8]. Consequently, efforts have concentrated on the development of robust frameworks to reliably detect deviations from GR using gravitational-wave signatures of compact-binary mergers.

One of these frameworks is the so-called *Test Infrastructure for General Relativity* (TIGER) [14, 13, 9]. TIGER operates by computing the odds ratio between GR and a test model in which one or more of the post-Newtonian coefficients are allowed to deviate from the value predicted by GR. The interested reader is referred to Refs. [14, 13, 9] for the details of the method and for analysis of its robustness against various potential systematic effects. To account for unmodelled effects, TIGER constructs a "background" distribution of odds ratios between GR and the test hypothesis by analysing $O(10^3)$ simulated GR signals. The background distribution defines the null hypothesis against which any particular observation (or catalog of observations) is tested. For validation purposes, the sensitivity of the algorithm to a specific deviation from GR is currently assessed by comparing it with a "foreground" odds ratio distribution. The foreground distribution is constructed by simulating a variety of signals in which the chosen deviation from GR is introduced. If the *integrated* overlap between the foreground and background distributions is smaller than a given false alarm probability, sensitivity to that particular deviation can be claimed.

The process described in the previous paragraph is extremely computationally expensive. If arbitrary combinations of k post-Newtonian coefficients are allowed to deviate from GR values, the total number of simulations that necessary to construct the background is 2^k for each synthetic source.

As a means to cheaply evaluate the detectability of particular deviations from GR, Cornish et al. [11] proposed an approximate formula to calculate the odds ratio between GR and an alternative model for gravity (AG). Subsequently, Vallisneri [15] proposed a similar approximation derived from the Fisher matrix formalism. Vallisneri's approximation considers the distribution of the odds ratio in the presence of noise and caracterises the efficiency and false alarm of a Bayesian detection scheme for alternative theories of gravity. Whilst neither of these approaches can replace the necessary analysis for real data, the possibility of having a quick and easily understandable formalism to check the performance of complex pipelines such as TIGER and assess whether a specific type of deviation is detectable without having to run thousands of simulations seems quite attractive.

In this work we investigate, in an idealised and controlled scenario, whether the predictions from Refs. [11, 15] are in agreement with the output of a numerical Bayesian odds-ratio calculation. We find in particular that the the analytical prescription of Ref. [15] is in reasonable agreement with the numerical result when

the fitting factor (FF) between AG and GR waveforms is ≥ 0.9, and that for FF ≤ 0.8, both analytical prescriptions overestimate the exact odds ratio.

Nevertheless, when the analytical odds ratio is regarded as an upper limit, useful indications of the detectability of a given deviation from GR can be drawn.

We analytically correct the approximate framework for computing the Bayes factor by introducing terms that are negligible at $FF \sim 1$, reproducing the proposed analytical expressions given in [11, 15], but become significant at lower values of the fitting factor. We show that these corrections extend the range of validity of the approximate expressions at least down to fitting factor values of ~ 0.7.

The rest of the paper is organised as follows: in section 2 we briefly review the Bayesian definition of the odds ratio; in section 2.1 we introduce the formula from Ref. [15]. In section 3 we present our findings and finally we discuss them in section 4.

2. Bayesian Inference for gravitational wave signals

In a Bayesian context, the relative probability of two or more alternative hypotheses given observed data d is described by the odds ratio (see, e.g., [10]). If GR is the general relativity hypothesis and AG is the hypothesis corresponding to some alternative theory of gravity, the odds ratio is given by:

$$O_{\rm AG,GR} = \frac{p({\rm AG}|d)}{p({\rm GR}|d)} = \frac{p({\rm AG})}{p({\rm GR})}\frac{p(d|{\rm AG})}{p(d|{\rm GR})} \equiv \frac{p({\rm AG})}{p({\rm GR})}B_{\rm AG,GR} \tag{1}$$

where we introduced the Bayes factor $B_{AG,GR}$, which is the ratio of the marginalised likelihoods (or evidences). The marginal likelihood is the expectation value of the likelihood of observing the data given the specific model H under consideration over of the prior probability distribution for all the model parameters θ:

$$p(d|H) \equiv Z = \int d\theta\, p(d|\theta, H)\, p(\theta|H)\,. \tag{2}$$

With the exception of a few idealised cases, the integral (2) is, in general, not tractable analytically. In gravitational-wave data analysis, the parameter space is at least 9-dimensional (for binaries with components that are assumed to have zero spin), and up to 15-dimensional for binaries with arbitrary component spins, and the integrand is a complex function of the data and the waveform model. For stationary Gaussian noise

$$p(d|\theta, H) \propto \exp[-(d - h(\theta)|d - h(\theta))/2]\,, \tag{3}$$

where $h(\theta) \equiv h(\theta|H)$ is the model waveform given parameters θ and we introduced the scalar product

$$(a|b) \equiv 2\int_0^\infty df\, \frac{a(f)b(f)^* + a(f)^*b(f)}{S(f)} \tag{4}$$

with the one-sided noise power spectral density $S(f)$. We analysed data from a single detector with a noise spectral density corresponding to the zero-detuning, high-power Advanced LIGO design configuration [25].

2.1. Analytical Approximation

Vallisneri [15] proposed an analytical approximation to the integral (2). He considered the following assumptions:

- linear signal approximation leading to a quadratic approximation of the log likelihood;
- only one additional dimension is necessary to describe the AG model;
- uniform prior distributions for all parameters describing both GR and AG models;
- the distance between the AG waveform and the manifold of GR waveforms is small so that the fitting factor (FF) between the two, defined as

$$FF = \left[\frac{(h_{\rm AG}|h_{\rm GR}(\theta))}{\sqrt{(h_{\rm AG}|h_{\rm AG})(h_{\rm GR}(\theta)|h_{\rm GR}(\theta))}} \right]_{\max \text{over } \theta} , \tag{5}$$

  is close to unity.

With the above assumptions, the integral Eq. (2) can be approximately computed analytically and the Bayes factor (1) is then given by:

$$B_{\rm AG,GR} \approx \sqrt{2\pi} \frac{\Delta\theta^a_{est}}{\Delta\theta^a_{prior}} e^{\rho^2(1-FF)} , \tag{6}$$

where ρ denotes the optimal signal-to-noise ratio (SNR)§:

$$\rho \equiv 2\sqrt{\int_0^\infty df \frac{|h(\theta_{\rm true})|^2}{S(f)}} . \tag{7}$$

The terms $\Delta\theta^a_{prior}$ and $\Delta\theta^a_{est}$ are the width of the prior distribution and of the Fisher matrix 1-σ uncertainty estimate for the additional AG parameter, respectively. Eq. (6) given here is valid for the case in which a zero realisation of the noise is present in the data. In Ref. [15] noise is considered and the appropriate formulae for the distribution of the Bayes factor over noise realisations can be found there. We opted for a zero-noise case for ease of comparison.

3. Comparison between the exact calculation and the analytical approximation

We compare the prediction from Eq. (6) with the evidence calculated by the Nested Sampling algorithm [18] as implemented in `lalinference` [17] in a simple experiment. Using the test waveform model presented in [14], we generate inspiral signals which would span a range of FFs. The testing waveform is a frequency-domain stationary phase approximation waveform, based on the TaylorF2 approximant [16], that has been modified in such way that the post-Newtonian coefficients are allowed to vary around the GR values within a given range. The TaylorF2 waveform for a face-on, overhead binary is given by:

$$h(f) = \frac{1}{D}\sqrt{\frac{5}{24}}\pi^{-2/3}\mathcal{M}^{5/6}f^{-7/6}e^{i\Psi(f)}, \tag{8}$$

§ Note that the definition by Vallisneri of the signal-to-noise ratio is different from ours. In Ref. [15] the signal-to-noise ratio quantity that appears in Eq. (6) is the norm of an hypothetical GR signal whose parameters are exactly the same as the "true" AG waveform, but with the extra AG parameter set to zero. In our case, the signal-to-noise ratio corresponds to the power in the AG signal (in the ideal case when it is filtered with AG templates). However, when the AG parameter is present only in the phase of the gravitational wave these two signal-to-noise ratios coincide.

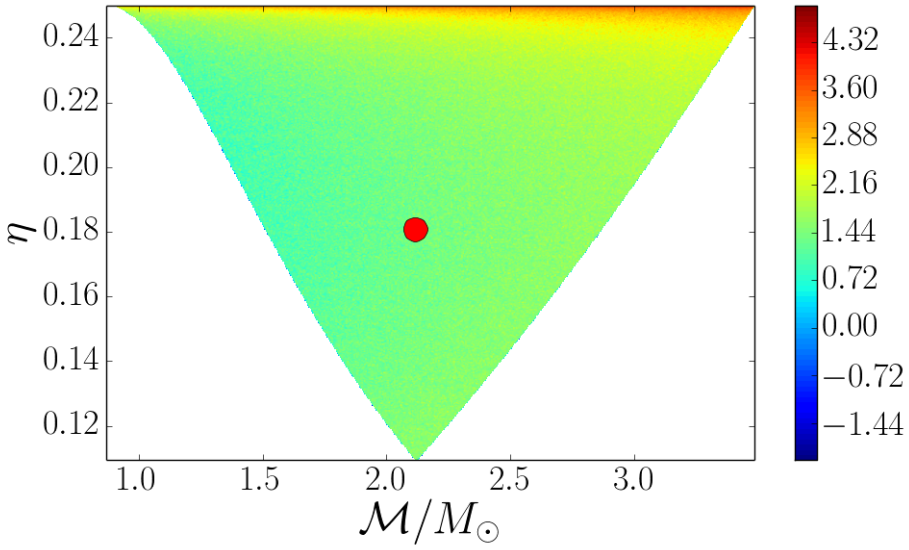

Figure 1. Two-dimensional logarithmic prior probability distribution in the $\mathcal{M}$ and η space. The shape of the region is set by the allowed ranges of individual and total masses. The red marker indicates the location of the $1.4M_\odot + 4.5M_\odot$ system simulated for the analysis.

where D is the luminosity distance, $\mathcal{M}$ is the chirp mass and the phase $\Psi(f)$ is:

$$\Psi(f) = 2\pi f t_c - \phi_c - \pi/4 + \sum_{i=0}^{7} \left[\psi_i + \psi_i^{(l)} \ln f\right] f^{(i-5)/3}. \tag{9}$$

The explicit forms of the coefficients ψ_i and $\psi_i^{(l)}$ in $(\mathcal{M}, \eta)$, where η is the symmetric mass ratio, can be found in [23]. In all our experiments we kept the parameters of the simulated sources fixed with the exception of the 1.5 post-Newtonian coefficient ψ_3 which we varied between $[0.5, 1.5]$ times its GR value by adding an arbitrary shift $d\chi_3$ between $[-0.5, 0.5]$:

$$\psi_3 \to \psi_3(1 + d\chi_3)\,. \tag{10}$$

The Nested Sampling algorithm was set up to sample from the following prior:

- the component masses where allowed to vary uniformly $\in [1, 7]M_\odot$ with the total mass constrained to the range $\in [2, 8]M_\odot$. This choice results in an allowed region of triangular shape in the $\mathcal{M}, \eta$ plane, see Fig. 1;
- uniform on the 2-sphere for sky position and orientation parameters;
- uniform in Euclidean volume for the luminosity distance;
- for recovery with AG templates, we used only one free testing parameter ($d\chi_3$) which was allowed to vary uniformly between $[-0.5, 0.5]$ times its GR value.

The FFs were computed from the maximum likelihood values obtained from the `lalinference` simulations, see Appendix A. The parameter uncertainty for Eq. (6) was computed using a 5-dimensional Fisher matrix calculation in which we varied

the two mass parameters, the time of coalescence, the phase at coalescence and the deviation parameter $d\chi_3$.

Our experiments were performed analysing simulated signals from a system whose component masses were chosen to be $1.4M_\odot + 4.5M_\odot$. We chose this system because it lies in the centre of our prior probability distribution over the masses, far away from prior boundaries. This minimises the impact of the prior on the fitting factor and Bayes factor computations, which ensures that we can make a fair comparison with Eq. (6) derived under the assumption of a uniform prior.‖

The approximate formula in Eq. (6) depends essentially on two quantities: the signal-to-noise ratio ρ and the FF. Below, we describe our investigations of the dependence of the Bayes factor on these two quantities.

For Nested Sampling calculations, the uncertainty on the calculated value of the evidence Z is evaluated as [18]:

$$\Delta \log Z \simeq \sqrt{\frac{H}{n}} \tag{11}$$

where H is the Kullback–Leibler divergence (or relative entropy) between the posterior distribution and the prior distribution and n is the number of live points used for Nested Sampling. H is computed by the Nested Sampling along side the evidence Z. Tipical values for $\Delta \log Z$ are $O(10^{-1})$.

Finally, it was recently pointed out that when a signal terminates abruptly in the detector band, measurement uncertainty may be significantly smaller than predicted by the Fisher matrix calculation [22]. To avoid these complications, we limited our analysis to frequencies between 30 and 512 Hz.

3.1. Scaling with the signal-to-noise ratio

We investigated the dependence of the Bayes factor on SNR by comparing the output of `lalinference` and Eq. (6) for a $1.4M_\odot + 4.5M_\odot$ system at SNRs of 10, 20, 30 and 40 at a fixed value of the FF, 0.992. Fig. 2 shows the Bayes factors from the two calculations.

The quadratic dependence of the Bayes factor on the SNR was verified by means of a simple chi-squared fit to an expression of the form

$$\ln B_{\mathrm{AG,GR}} = \alpha \mathrm{SNR}^\beta + \gamma\,. \tag{12}$$

The scaling of $\log(B_{\mathrm{AG,GR}})$ with the SNR of the appears to be consistent with the expected quadratic dependence: we find $\beta = 1.95 \pm 0.4$.

3.2. Scaling with the fitting factor

We evaluate the dependence of the Bayes factor on FF by again injecting a signal from a $1.4M_\odot + 4.5M_\odot$ binary, now at a fixed SNR of 20 but with varying FF. As in the previous section, we vary the FF by adding arbitrary deviations from the GR value to the 1.5 post-Newtonian phase coefficient. In particular, we varied $d\chi_3$ between -0.5 and 0.5, leading to FF $\in [0.7, 1.0]$. We verified that our injection was sufficiently far from prior boundaries by confirming that the Bayes factor is the same for positive and negative values of $d\chi_3$ that yield the same FF.

‖ For example, an equal mass system would lie exactly on the prior boundary at $\eta = 0.25$. For this reason, the GR model has very little room in the η direction to accommodate the additional phase shift due to a non-zero $d\chi_3$. The net result is a very rapid drop in FF towards negative $d\chi_3$.

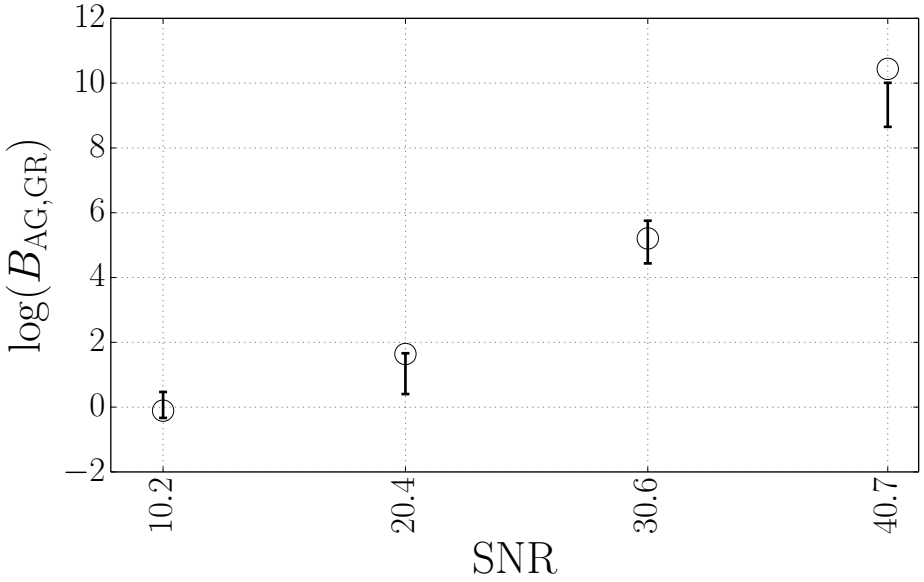

Figure 2. Logarithmic Bayes factor from `lalinference` (errorbars) and from Eq. (6) (empty circles) as a function of the SNR, for FF fixed to 0.992. The `lalinference` errors are computed from Eq. (11). The circles are the values of the logarithmic Bayes factor obtained from Eq. (6) using the FF extracted from the maximum likelihood values as computed by `lalinference`.

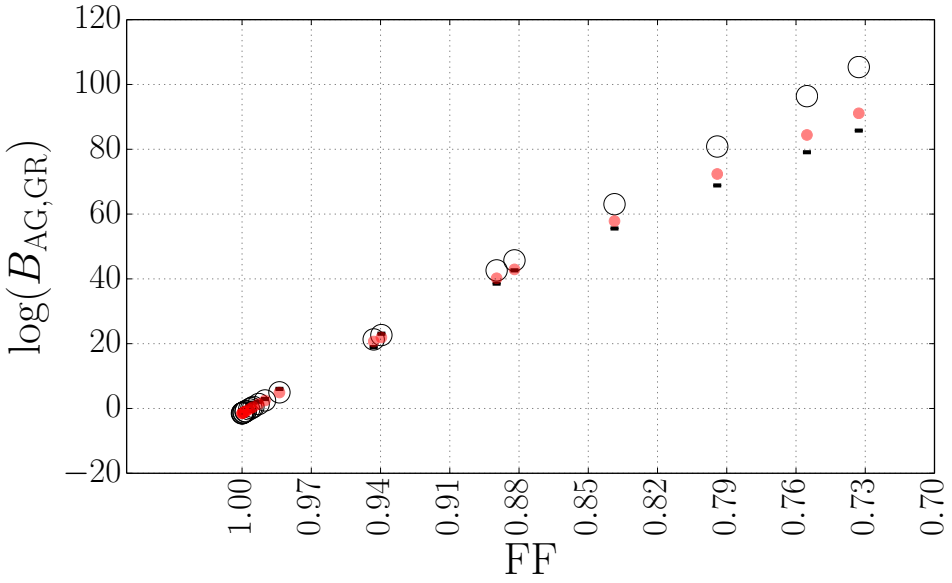

Figure 3. Logarithmic Bayes factor from `lalinference` (error bars) and from Eq. (6) (empty circles) as a function of the FF. The SNR was fixed to 20. The Bayes factors computed with the two approaches agree for FF$\sim$ 1 but tend to diverge for decreasing FF. The red dots indicate the value of the logarithmic Bayes factor obtained by using a quadratic dependence on the FF, Eq. (16), rather than the linear dependence of Eq. (6).

Fig. 3 shows the logarithmic Bayes factor computed by `lalinference` and from Eq. (6). The two methods agree for FF ~ 1. At FF ≤ 0.9, the analytical approximation overestimates the value of the Bayes factor compared to `lalinference`. Moreover, the disagreement gets worse with decreasing FF, suggesting a nonlinear dependence on the FF. In the next section, we investigate the approximate analytical expression in greater detail and derive additional corrections that extend its validity to lower fitting factors.

3.3. Correcting the analytical expression for lower fitting factors

Under the assumption that the region of likelihood support on the parameter space is small, and that over this region the prior does not vary significantly, the evidence for any of the models H_i, depending on parameters θ, under consideration can be approximated as (e.g., [24], correcting for a typo in the exponent of (2π)):

$$Z(H_i) \propto [L_{H_i}]_{\max \text{ over } \theta} \, (2\pi)^{N/2} \prod_i^N \frac{\Delta\theta^i_{est}}{\Delta\theta^i_{prior}} \,, \tag{13}$$

where N is the number of parameters. Strictly speaking, the equation above is only valid when parameters are uncorrelated. In the general case of correlated parameters, $\prod_i^N \Delta\theta^i_{est}$ should be replaced with the uncertainty volume in which the likelihood has support, while $\prod_i^N \Delta\theta^i_{est}$ is shorthand for the total prior volume. However, such correlations do not affect the scaling of the uncertainty with the SNR, and do not impact our conclusions.

Therefore, the Bayes factor between the AG and GR model can be approximated as the ratio of the maximum likelihoods times the product of the ratios of posterior widths to prior supports:

$$B_{AG,GR} \approx \frac{[L_{AG}]_{\max \text{ over } \theta'}}{[L_{GR}]_{\max \text{ over } \theta}} \sqrt{2\pi} \frac{\prod_{i=0}^{N} \frac{\Delta\theta'_{est}}{\Delta\theta'_{prior}}}{\prod_{i=0}^{N-1} \frac{\Delta\theta_{est}}{\Delta\theta_{prior}}} \tag{14}$$

where θ' and θ are parameter vectors within the AG and GR models, respectively, and N is the dimensionality of the AG parameter space.

We begin by considering just the first term in Eq. (14), which scales exponentially with the SNR in contrast to the components of the second term, which scale inversely with the SNR. Neglecting the second term, we find:

$$\log(B_{AG,GR}) \propto \log([L_{AG}]_{\max \text{ over } \theta'}) - \log([L_{GR}]_{\max \text{ over } \theta}) \,. \tag{15}$$

Using Eq. (A.9), we find:

$$\log(B_{AG,GR}) \propto \frac{\rho^2}{2}(1 - FF^2) \,. \tag{16}$$

which is the expression originally proposed in [11]. At FF close to unity, $(1 - \mathrm{FF}^2) = (1 + \mathrm{FF})(1 - \mathrm{FF}) \approx 2(1 - \mathrm{FF})$, the approximation implicitly made in [15], and we recover Eq. (6). However, we expect (16) to lead to a better fit at low fitting factors. The filled (red) dots in Fig. 3 show the Bayes factor computed via Eq. (16), with the proportionality constant fixed to be the same as in Eq. (6). Indeed, Eq. (16) predicts Bayes factors that are in closer agreement with the exact ones than Eq. (6). In this case, disagreements with the exact result can be seen for FF ~ 0.75, when the differences in the local shapes of the GR and AG manifolds can become significant.

Vallisneri [15] further assumed that the priors and measurement uncertainties on all parameters except the one describing the deviation from GR, θ^a, are the same for the AG and GR models (which, in turn, is a statement about the similarity in the shape of the two waveform manifolds near the maximum likelihood locations). In this case, the Bayes factor between the two models is [cf. (6)]:

$$B_{AG,GR} \propto \frac{[L_{AG}]_{\max \text{ over } \theta'}}{[L_{GR}]_{\max \text{ over } \theta}} \sqrt{2\pi} \frac{\Delta\theta^a_{est}}{\Delta\theta^a_{prior}} \,, \tag{17}$$

where ^{a} again refers to the one additional AG parameter which describes the deviation from GR.

However, we should not expect that the posterior widths will be identical in the AG and GR models for all parameters except the additional AG parameter are the same in the AG and GR models. At high SNRs where the log likelihood can be approximated by a quadratic, posterior widths should scale inverse with the signal-to-noise ratio ρ. While ρ is the optimal SNR recovered when AG templates are used within the AG model, the maximal SNR recoverable when using GR templates within the GR model is lower. By definition, this GR SNR is

$$\rho_{\rm GR} \equiv \left[\frac{(h_{\rm AG}|h_{\rm GR}(\theta))}{\sqrt{(h_{\rm GR}(\theta)|h_{\rm GR}(\theta))}} \right]_{\max \text{ over } \theta} = {\rm FF}\, \rho \,. \tag{18}$$

Assuming the inverse SNR scaling of the posteriors, and using identical priors on common parameters in the AG and GR models, Eq. (14) reduces to

$$B_{\rm AG,GR} \approx \frac{[L_{\rm AG}]_{\max \text{ over } \theta'}}{[L_{\rm GR}]_{\max \text{ over } \theta}} {\rm FF}^{N-1} \sqrt{2\pi} \frac{\Delta\theta^a_{est}}{\Delta\theta^a_{prior}} \,. \tag{19}$$

Taking a logarithm of this equation and again using $\frac{\rho^2}{2}(1 - {\rm FF}^2)$ for the difference between maximum likelihoods (A.9), we find

$$\begin{aligned} \log(B_{\rm AG,GR}) \approx & \frac{\rho^2}{2}(1 - {\rm FF}^2) + (N-1)\log({\rm FF}) \\ & + \log\left(\sqrt{2\pi} \frac{\Delta\theta^a_{est}}{\Delta\theta^a_{prior}} \right) . \end{aligned} \tag{20}$$

Eq. (20) reduces to Eq. (6) for ${\rm FF} \sim 1$. However, it is accurate for a much wider range of fitting factors. Fig. 4 shows the comparison between the log-Bayes factors from `lalinference` (error bars), the ones from Eq. (6) (circles) and finally the ones from Eq. (20) (red dots). Indeed, the log-Bayes factors from Eq. (20) show a very close agreement with the numerical values. Thus, Eq. (20) provides a good approximation to the exact values of the log-Bayes factors.

Another merit of Eq. (20) is that it can be generalised to an arbitrary number of extra non-GR parameters. If we have k non-GR parameters, Eq. (20) becomes:

$$\begin{aligned} \log(B_{\rm AG,GR}) \approx & \frac{\rho^2}{2}(1 - {\rm FF}^2) + (N-k)\log({\rm FF}) \\ & + \log\left((2\pi)^{k/2} \prod_i^k \frac{\Delta\theta^{a_i}_{est}}{\Delta\theta^{a_i}_{prior}} \right) . \end{aligned} \tag{21}$$

Throughout this work, we have restricted our attention to the zero-noise realisation. Vallisneri analysed the distribution of the Bayes factor under different noise realisations and showed (see Eq. (15) of [15]) that fluctuations in the logarithm

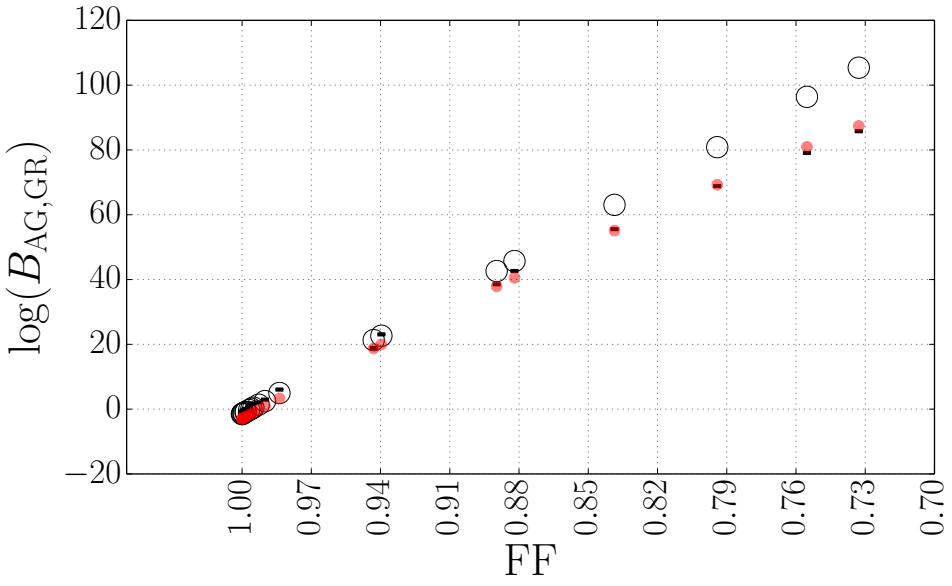

Figure 4. Logarithmic Bayes factor from `lalinference` (error bars), from Eq. (6) (empty circles), and from Eq. (20) (red dots) as a function of the FF. The SNR is fixed to 20. The `lalinference` errors are computed from Eq. (11). The extended analytical expression of Eq. (20) which includes corrections for lower fitting factors is in good agreement with the exact calculation from `lalinference`.

of the Bayes factor have a standard deviation of $\sim \sqrt{2}\rho\sqrt{1-\mathrm{FF}}$. While our additional corrections to the Bayes factor also lead to corrections in this quantity, we neglect these second-order effects.

We can compare the two systematic corrections discussed above to the level of these statistical fluctuations due to noise. The difference in the log-Bayes factor between Eq. (6) and Eq. (16), i.e., the difference between the approximations of Refs. [15] and [11], is $(1/2)\rho^2(1-\mathrm{FF}^2) - \rho^2(1-\mathrm{FF}) = -(1/2)\rho^2(1-\mathrm{FF})^2$. This difference is approximately equal to the statistical fluctuation in the log-Bayes factor for $\rho = 20$ and $\mathrm{FF} \sim 0.73$, corresponding to the rightmost points in Fig. 4. Meanwhile, the new correction to the log-Bayes factor which we introduced in Eq. (20) has a magnitude of $(N-1)\log \mathrm{FF}$; for $N = 9$ and other parameters as above, it is several times smaller than the noise-induced fluctuations.

Therefore, these corrections are unlikely to impact the detectability of a deviation from GR; in any case, in practice the detectability of the deviation would be determined by an analysis of the data and a numerical computation of the Bayes factor, not approximate predictive techniques. However, these corrections are useful in explaining the apparent difference between numerical and analytical calculations, and therefore help validate both approaches by enabling a successful cross-check.

4. Discussion

We computed the Bayes factor between a GR model and an alternative gravity model for a gravitational-wave signature of an inspiraling compact binary. We compared two

calculations of the Bayes factor: an exact numerical computation with `lalinference` and an approximate analytical prediction due to Vallisneri [15]. We verified that the analytical approximation yields the correct scaling of the logarithm of the Bayes factor with the square of the signal-to-noise ratio at high fitting factor values. However, the predicted scaling of the Bayes factor with the fitting factor is inaccurate for FF ≤ 0.9.

We extended the regime of validity of the analytical approximation of [15] to lower fitting factors by including additional FF-dependent terms and by extending to multiple non-GR parameters. We confirmed that the more complete analytical prediction that we derived in this work, Eq. (20), remains valid down to fitting factors of ≤ 0.7.

It is worth noting that Eq. (6) loses accuracy precisely in the regime where it becomes possible to differentiate GR and alternative gravity models. The FF is very close to unity in the regime in which the GR waveform can still match a signal which violates GR through different choices of the values of the binary's parameters within the GR framework. The Bayes factor in this case is not significantly different from 1, thus no decision on the nature of the signal can be made at an acceptably low false alarm probability¶. Therefore, our extension of the analytical expression for the Bayes factor to lower fitting factors provides a useful, easy-to-compute approximate technique precisely in the regime of interest in the case of a zero noise realisation.

The analytical expressions presented in [11, 15] and in this work are predicated on the assumption that the $(N-1)$–dimensional GR manifold and the N–dimensional AG manifold are sufficiently similar near the maximum-likelihood values that the parameter uncertainties can be assumed to be equal (up to scaling with the inverse SNR) on the two manifolds. Differences in the local curvature of the two manifolds could become significant when the distance between them is large, or the systematic bias between true and best-fit parameters is significant relative to statistical measurement uncertainty. Therefore, this assumption could (although need not) break down either at small fitting factors, or, somewhat paradoxically, at large SNR for a fixed fitting factor. Specifically, when $\rho^2(1-FF^2) \gg N$, the uncertainty region within a manifold is much smaller than the distance between manifolds or between true and best-fit parameters within a manifold, and the AG and GR manifolds may no longer yield similar parameter uncertainties.

Another possible cause of the breakdown of the analytical approximation is the impact of priors. If the prior distribution is very non-uniform within the region of likelihood support, particularly if a sharp prior boundary is present within this region, the analytical approach described above is no longer valid. A further limitation is the restriction to high SNRs. The widths of the posterior distributions are inversely proportional to the SNR only when the linearized-signal approximation is valid (i.e., when the covariance matrix is well approximated by the inverse of the Fisher matrix).

In summary, the analytical approximation presented by Cornish et al. [11] and Vallisneri [15], and its extensions as given in Eqs. (20) and (21), provide a computationally cheap way of predicting the detectability of a deviation from GR for a given AG theory without the need to run expensive numerical simulations, subject to the limitations outlined above. Hence, these analytical approximations can be a very useful tool to get quick indications of whether a particular class and magnitude of one-parameter deviations from GR are detectable. However, these methods are

¶ This regime is a case of the so-called "fundamental bias" [19]. It is treated using the analytical approximation presented in [15] by [20]. A numerical study with the `lalinference` code can be found in [21].

merely predictive, and inference on actual data must rely on parameter estimation and model comparison with complete data-analysis pipelines.

Acknowledgements

We thank Michele Vallisneri, Neil Cornish, John Veitch, Will Farr, Christopher Berry, Carl-Johan Haster and Zachary Hafen for useful comments and discussions. The work was funded in part by a Leverhulme Trust research project grant. The numerical simulations were performed on the Tsunami cluster of the University of Birmingham.

[1] B. Abbott et al., Rept. Prog. Phys. **72** 076901 (2009) [gr-qc/0711:3041]
[2] G.M. Harry for the LIGO Scientific Collaboration, *Class. Quant. Grav.* **27**, 084006 (2010)
[3] F. Acernese et al., Classical and Quantum Gravity **25**, 184001 (2008)
[4] T. Accadia et al., Classical and Quantum Gravity **28**, 114002 (2011)
[5] T. Accadia et al., Journal of Instrumentation, **7** P03012 (2012)
[6] `\protect\vrule width0pt\protect\href{http://wwwcascina.virgo.infn.it/advirgo}{http://wwwcascina.virgo.infn.it/advirgo}`
[7] Abadie, J., Abbott, B. P., Abbott, R., et al. 2010, Classical and Quantum Gravity, 27, 173001
[8] Burgay, M., D'Amico, N., Possenti, A., et al. 2003, Nature, 426, 531
[9] Agathos, M., Del Pozzo, W., Li, T. G. F., et al. 2014, *Phys. Rev.* D, 89, 082001
[10] Del Pozzo, W., Veitch, J., & Vecchio, A. 2011, *Phys. Rev.* D, 83, 082002
[11] Cornish, N., Sampson, L., Yunes, N., & Pretorius, F. 2011, *Phys. Rev.* D, 84, 062003
[12] Sampson, L., Cornish, N., & Yunes, N. 2013, *Phys. Rev.* D, 87, 102001
[13] Li, T. G. F., Del Pozzo, W., Vitale, S., et al. 2012, Journal of Physics Conference Series, 363, 012028
[14] Li, T. G. F., Del Pozzo, W., Vitale, S., et al. 2012, *Phys. Rev.* D, 85, 082003
[15] Vallisneri, M. 2012, *Phys. Rev.* D, 86, 082001
[16] Buonanno, A., Iyer, B. R., Ochsner, E., Pan, Y., & Sathyaprakash, B. S. 2009, *Phys. Rev.* D, 80, 084043
[17] Farr, B. et al, in preparation
[18] Skilling, J. 2004, American Institute of Physics Conference Series, 735, 395
[19] Yunes, N., & Pretorius, F. 2009, *Phys. Rev.* D, 80, 122003
[20] Vallisneri, M., & Yunes, N. 2013, *Phys. Rev.* D, 87, 102002
[21] Vitale, S., & Del Pozzo, W. 2014, *Phys. Rev.* D, 89, 022002
[22] Mandel, I., Berry, C., Ohme, F., Fairhurst, S., Farr., W. 2014, arXiv:1404.2382
[23] C.K. Mishra, K.G. Arun, B.R. Iyer, and B.S. Sathyaprakash, Phys. Rev. D **82**, 064010 (2010); arXiv:1005.0304
[24] Veitch, J., Mandel, I., Aylott, B., Farr, B., Raymond, V., Rodriguez, C., van der Sluys, M., Kalogera, V., Vecchio, A. 2012, *Phys. Rev.* D 85, 104045
[25] *Advanced LIGO anticipated sensitivity curves*, `https://dcc.ligo.org/cgi-bin/DocDB/ShowDocument?docid=2974`
[26] Schwarz, G.E., 1978, Annals of Statistics 6 (2): 461464, MR 468014

Appendix A. Computing Fitting Factors from logLikelihoods

The fitting factor Eq. (5) can be extracted directly from the Nested Sampling runs without the need to search over a parameter grid. Begin by writing the logarithmic likelihood in a zero noise realisation:

$$\log(L) = const + (h_{true}|h(\theta)) - \frac{(h_{true}|h_{true})}{2} - \frac{(h(\theta)|h(\theta))}{2} \qquad \text{(A.1)}$$

where h_{true} is the gravitational wave signal in the data stream, $h(\theta)$ is the search template, and *const* is a constant. Consider the difference $\Delta\lambda$ between the maximum log likelihoods given for the AG and GR models given an AG signal:

$$\Delta\lambda = \log(L_{GR})_{\max \text{over } \theta} - \log(L_{AG})_{\max \text{over } \theta'}\,, \qquad \text{(A.2)}$$

where

$$\log(L_{GR}) = const + (h_{true}|h_{GR}(\theta)) - \frac{(h_{true}|h_{true})}{2} - \frac{(h_{GR}(\theta)|h_{GR}(\theta))}{2} \tag{A.3}$$

and $\log(L_{AG})_{\max \text{over } \theta'} = const$ since this likelihood is maximized for $h_{AG}(\theta') = h_{true}$.

We can maximise the GR log-likelihood analytically over the amplitude of $h_{GR}(\theta)$ by defining

$$h_{GR}(\theta) = A\hat{h}_{GR}(\xi) \tag{A.4}$$

with $\xi \equiv \theta \setminus A$ being the set of parameters other than the amplitude. One can solve for the value of the amplitude that satisfies

$$\frac{\partial \log(L_{AG})}{\partial A} = 0 \tag{A.5}$$

and obtain:

$$A = \frac{(h_{true}|\hat{h}_{GR}(\xi))}{(\hat{h}_{GR}(\xi)|\hat{h}_{GR}(\xi))}\,. \tag{A.6}$$

Substituting this into Eq. (A.3) and setting $(h_{true}|h_{true}) \equiv \rho^2$ yields:

$$\Delta\lambda = \frac{1}{2}[\frac{(h_{true}|\hat{h}_{GR}(\xi))^2}{(\hat{h}_{GR}(\xi)|\hat{h}_{GR}(\xi))}]_{\max \text{over } \xi} - \frac{\rho^2}{2}\,. \tag{A.7}$$

Meanwhile, the fitting factor FF can be similarly written as:

$$FF = [\frac{(h_{true}|\hat{h}_{GR}(\xi))}{\rho\sqrt{(\hat{h}_{GR}(\xi)|\hat{h}_{GR}(\xi))}}]_{\max \text{over } \xi}\,. \tag{A.8}$$

Thus,

$$\Delta\lambda = \frac{1}{2}\rho^2 FF^2 - \frac{\rho^2}{2} = -\frac{\rho^2}{2}(1 - FF^2) \tag{A.9}$$

and

$$FF = \sqrt{\frac{2\Delta\lambda}{\rho^2} + 1}\,. \tag{A.10}$$

The numerically computed values of $\Delta\lambda$ have an intrinsic variability due to the stochastic nature of the sampler. The standard deviation for $\Delta\lambda$ derived from our simulations is $\sigma_{\Delta\lambda} = 0.016$. The corresponding uncertainty in our FF estimate is given by

$$\sigma_{FF} = \frac{1}{\rho^2 FF}\sigma_{\Delta\lambda}\,. \tag{A.11}$$

For an SNR of 20 and a FF of 1, $\sigma_{FF} = 4 \times 10^{-5}$. For all practical purposes we can consider our estimated FFs to be exact.